

Sparse convolutional coding for neuronal ensemble identification

Sven Peter¹, Daniel Durstewitz², Ferran Diego¹, and Fred A. Hamprecht¹

¹Heidelberg Collaboratory for Image Processing, Heidelberg University

²Department of Theoretical Neuroscience, Bernstein Center for Computational Neuroscience, Central Institute of Mental Health, Medical Faculty Mannheim of Heidelberg University

June 23, 2016

Cell ensembles, originally proposed by Donald Hebb in 1949, are subsets of synchronously firing neurons and proposed to explain basic firing behavior in the brain. Despite having been studied for many years no conclusive evidence has been presented yet for their existence and involvement in information processing such that their identification is still a topic of modern research, especially since simultaneous recordings of large neuronal population have become possible in the past three decades. These large recordings pose a challenge for methods allowing to identify individual neurons forming cell ensembles and their time course of activity inside the vast amounts of spikes recorded. Related work so far focused on the identification of purely simultaneously firing neurons using techniques such as Principal Component Analysis. In this paper we propose a new algorithm based on sparse convolution coding which is also able to find ensembles with temporal structure. Application of our algorithm to synthetically generated datasets shows that it outperforms previous work and is able to accurately identify temporal cell ensembles even when those contain overlapping neurons or when strong background noise is present.

1 Introduction

Cell ensembles (or synonymously *cell assemblies* or *cortical motifs*) were originally proposed by Hebb [1] as subsets of synchronously firing neurons to explain brain activity underlying complex behaviors. Multiple studies have been done to find evidence for or against the neuronal ensemble hypothesis on datasets recorded from different areas of the brain inside different animals but no clear answer has been found yet [5, 3, 18, 9, 17].

Especially in past two decades where it became possible to record large neuronal populations concurrently [13, 24, 25] methods such as Principal Component Analysis (PCA) [11] / Singular Value Decomposition (SVD) [32], Independent Component Analysis (ICA) [29] or Non Negative Matrix Factorization (NMF) [26] have been applied to identify neurons repeatedly firing at the same time (see figure 1.1a) to find statistically significant ensembles and answer the question about their existence. While there are also more complex methods dealing with jitter in individual spike times [30] or the recurrence of motifs involving the ensembles themselves [26, 32] more complex motifs such as synfire chains [14] (see figure 1.1b) or motifs where only a single neuron is active at a certain time (see figure 1.1c) are still missed [29].

With the increasing size of available datasets more powerful methods are required to not only handle low signal-to-noise levels but to also understand the formation and existence of ensembles with more complex temporal structures. In this paper we leverage sparsity constraints on neuronal activity to allow a simple and elegant mathematical formulation to identify ensembles completely unsupervised. We show that our proposed algorithm outperforms state-of-the-art methods on synthetic data.




Figure 1.1: Synchronous and temporal motifs: All three illustrations show four neurons spiking at various times. The spikes highlighted in red are part of a repeating motif. In (a) the pattern is just the synchronous activity of all neurons while the synfire chain in (b) and the motif in (c) also contain temporal structure, respectively

2 Related work

Santos, Ribeiro, and Tort [29] provide an overview of using principal and independent component analysis (PCA and ICA) to identify ensembles. Both require an estimation of the number of ensembles N_e at first which is based on the number of significant eigenvalues of the correlation matrix. Significance is either established using random matrix eigenvalue distribution theory [2] or by shuffling the spike matrix to remove temporal correlation while preserving spike count distribution.

PCA, which has been used for a long time to track cell ensembles [7], computes the first N_e principal components of the spike matrix and considers those to be the ensembles. Its biggest limitations are that two different ensemble patterns can be merged into a single component and that neurons shared between two ensembles are assigned lower weights than expected. Additionally negative values with no physical meaning are possible in the components. Recovering individual neurons which belong to a single ensemble is not reliably possible [23, 29].

ICA decomposes a multivariate signal into additive subcomponents with the assumption that these are non-Gaussian and statistically independent from each other [6]. When used to learn ensembles it overcomes some of the problems of PCA-based methods: Individual neuron-ensemble membership can be recovered easily and neurons belonging to multiple ensembles are also correctly identified [29]. Again negative values are possible in the identified patterns leading to interpretation problems. For synchronous patterns Santos, Ribeiro, and Tort [29] recommend to use this method since it provides the best estimate. Temporal structure is however not part of this model and patterns such as synfire chains cannot be identified.

Carrillo-Reid et al. [32] use singular value decomposition (SVD) to identify synchronously firing neurons. They identify directionally sensitive ensembles in the visual cortex of mice and are able to correlate their activity to the time during which external stimuli were presented. They are also able to identify ensembles sensitive to natural scenes and show repeating activations of those by applying graph theoretical methods to the already identified ensembles. Unlike our method however the number of ensembles has to be estimated at first and only synchronously spiking neurons are considered in the first step. While they can find activity spread over time these have to have at least two neurons firing synchronously in every bin, such that synfire chains would not be identified again.

Billeh et al. [30] tried to identify *almost* synchronously firing neurons: Instead of relying on almost perfect synchronicity or on appropriate binning to ensure this condition they take into account that individual neurons belonging to a single ensemble will show a small jitter in their spiking time and used a dynamics-driven methodology based on the Markov Stability framework [20] to identify ensembles at multiple levels of granularity. While they are able to identify connections between neurons even if the activity is shifted in time they do not identify the exact motifs: Their result is the connectivity between neurons while we identify the connections and their temporal relation simultaneously instead.

Diego and Hamprecht [26] propose to use non-negative matrix factorization techniques to decompose binned spike matrix into multiple levels of synchronous patterns to identify a hierarchical structure of motifs. Again no temporal structure is taken into account and only neurons firing at the same time are considered.

Approaches based on one dimensional convolutional coding [21] have been used to recover spike trains of individual neurons from recorded Calcium fluorescence sequences [31, 28, 27, 22]. In these models each neuron is however treated independently and they have not been extended to model relationships between neurons. In order to extract motifs a novel optimization approach is required to replace the one dimensional with two dimensional filters.

We propose to adapt the idea behind non negative convolutional matrix factorization [16], which has originally been developed to allow the extraction of motifs in audio processing, for the learning of motifs. This method only regularizes the activity of the learned motifs with a ℓ_1 prior which is too weak to recover motifs in neuronal spike trains. Instead we propose a different optimization technique that allows to regularize the motifs with a ℓ_1 and their activities with a much stronger ℓ_0 prior. This allows for a simple and elegant formulation to learn complex motifs from recorded spike trains.

3 Our method




Figure 3.1: Convolutional coding: In this example the raw signal \mathbf{Y} is an additive mixture of two motifs \mathbf{a}_0 and \mathbf{a}_1 (which have been highlighted with different colors) convolved with their activities \mathbf{s}_0 and \mathbf{s}_1 resulting in the reconstruction \mathbf{X} . Additional background noise \mathbf{N} consists of non-repeating spikes. Both the motifs and their activities are learned simultaneously.

Let $\mathbf{Y} \in \mathbb{R}_+^{n \times m}$ be a matrix whose n rows represent individual neurons with their spiking activity binned to m columns. We assume that this raw signal is an additive mixture of l of motifs $\mathbf{a}_i \in \mathbb{R}_+^{n \times \tau}$ with temporal length τ convolved with a sparse activity signal \mathbf{s}_i plus noise (see figure 3.1).

We address the unsupervised problem of simultaneously estimating both the coefficients making up the motifs \mathbf{a}_i and their activities \mathbf{s}_i . To this end, we propose to solve the following optimization problem:




$$\min_{\mathbf{a}, \mathbf{s}} \left\| \mathbf{Y} - \sum_i^l \mathbf{s}_i \circledast \mathbf{a}_i \right\|_F^2 + \alpha \sum_i^l \|\mathbf{s}_i\|_0 + \beta \sum_i^l \|\mathbf{a}_i\|_1 \quad (3.1)$$

The ℓ_0 norm is chosen for \mathbf{s} since Diego and Hamprecht [31] have successfully used it

to learn spike trains of neurons. For the ensembles themselves the ℓ_1 norm is used to enforce only few non-zero coefficients [15].

This problem is non-convex in general but can be approximated by initializing \mathbf{s} to random noise and using a block coordinate descent strategy [10, Section 2.7] to alternately optimize for the two variables. The activities \mathbf{s} are inferred using convolution matching pursuit [21, 4, 19] and the ensembles themselves using LASSO regression with non-negativity constraints [8] by transforming the convolution with \mathbf{s}_i to a linear set of equations using Toeplitz matrices $\tilde{\mathbf{s}}_i$ with $\tilde{s}_{i,j,k} = \tilde{s}_{i,j+1,k+1} = s_{i,j-k}$ for $j \geq k$ and $\tilde{s}_{i,j,k} = 0$ for $j < k$ (where i denotes the i th matrix with element indices j and k) which are then stacked next to each other: [12, 15]:

$$\begin{aligned} \min_{\mathbf{a}} & \left\| \text{vec}(\mathbf{Y}) - \sum_i^l \tilde{\mathbf{s}}_i \cdot \text{vec}(\mathbf{a}_i) \right\|_2^2 + \beta \sum_i^l \|\mathbf{a}_i\|_1 \\ &= \min_{\mathbf{a}} \left\| \underbrace{\text{vec}(\mathbf{Y})}_{b \in \mathbb{R}^{mn}} - \underbrace{[\tilde{\mathbf{s}}_0 \dots \tilde{\mathbf{s}}_l]}_{A \in \mathbb{R}^{mn \times ln\tau}} \underbrace{\begin{bmatrix} \text{vec}(\mathbf{a}_0) \\ \dots \\ \text{vec}(\mathbf{a}_l) \end{bmatrix}}_{x \in \mathbb{R}^{ln\tau}} \right\|_2^2 + \beta \sum_i^l \|\mathbf{a}_i\|_1 \end{aligned} \quad (3.2)$$



(a) Ground truth ensemble (b) Possible wrong ensembles after a single iteration (c) Equivalent padded ensembles

Figure 3.2: Ground truth ensemble and learned state after a single iteration: This figure shows four neurons forming a motif over three bins. After a single iteration parts of the motif can be missing which is solved by increasing the ensemble length and centering the motifs after each iteration.

Special care has to be taken to avoid missing parts of the motif due to the originally identified positions. Consider the ground truth ensemble seen in figure 3.2a. After a single iteration the learned ensemble could be any of the two wrong possibilities seen in figure 3.2b. While the learned motif does indeed occur in real data it is not complete and can never be completed since there is no more space on the left or right to identify the missing associations. To overcome this problem the vectors \mathbf{a} have to be chosen larger than required and centered after each iteration when possible: When there are enough empty rows on either side the whole motif can be shifted before the new ensemble activities \mathbf{s} are identified. This does not increase the reconstruction error since the activities will also just be shifted by the same amount. When new coefficients \mathbf{a} are learned in the next

iteration there now is enough space to also capture the previously missed associations (see figure 3.2c which can be completed in the next iteration).

4 Results on synthetic data

Since ground truth datasets are in general not available different synthetic data sets consisting of fifty neurons observed over one thousand time bins have been simulated to compare our method to existing work. A subset of the neurons is randomly assigned to belong to a single ensembles, others to multiple ensembles and the rest are not part of any ensemble and fire completely on their own. The ensemble activity itself is modeled as a Poisson process with a randomly chosen mean [29] and a refractory period of at least the length of the ensemble itself. Additionally spurious spikes of single neurons are added to simulate neurons firing out of sync. The percentile of neurons belonging to multiple ensembles, the fraction of spurious spikes and the temporal lengths of the ensembles have been varied to create different test cases.

For PCA and ICA based methods the number of ensembles is estimated using the Marchenko-Pastur eigenvalue distribution [29]. The sparsity parameter in O’Grady’s sparse convolutive non-negative matrix factorization that resulted in the best performance was experimentally chosen [16]. Additionally the normalized correlation between the spike trains of each possible neuron-neuron combination has been calculated. If the correlation between neuron i and neuron j is higher than the correlation between i and any other neuron those two are assumed to be connected within an ensemble.

A simplified example dataset with twenty neurons, three temporal motifs and fifty spurious spikes per neuron can be seen in figure 4.1. When running our method with two different random initial states to identify five motifs with a temporal length of fifteen frames for ten iterations all three original motifs are learned successfully (figure 4.1c and 4.1d; the motifs have been sorted manually to match up with the ground truth). The two spurious motifs change depending on the random initial state while the true motifs always appear in the results. Neither PCA, ICA or scNMF are able to learn any true motif.

To evaluate the different methods a neuron association matrix is calculated from the learned ensembles by choosing a threshold above which neurons are assumed to belong to an ensemble and compared to the ground truth association matrix. The functional association between neurons has been used as an indicator of performance in previous work [23, 32, 30] but does not take the identification of the correct temporal motifs into account. We still chose this method since it works without limitations for synchronous motifs and also allows comparisons for the more complex cases.

We plotted three different ROC curves for the different temporal lengths $\tau = 1, 7$ and 21 in figure 4.2. In the synchronous case (i.e. $\tau = 1$, figure 4.2a) our proposed method performs as good as the best competitor. As expected PCA performance shows a huge variance since some of the datasets contain neurons shared between multiple ensembles and since extracting actual neuron-ensemble assignments is not always possible [23, 29].

When temporal structure is introduced we are still able to identify associations between neurons with very high accuracy. Sparse convolutional matrix factorization is able to

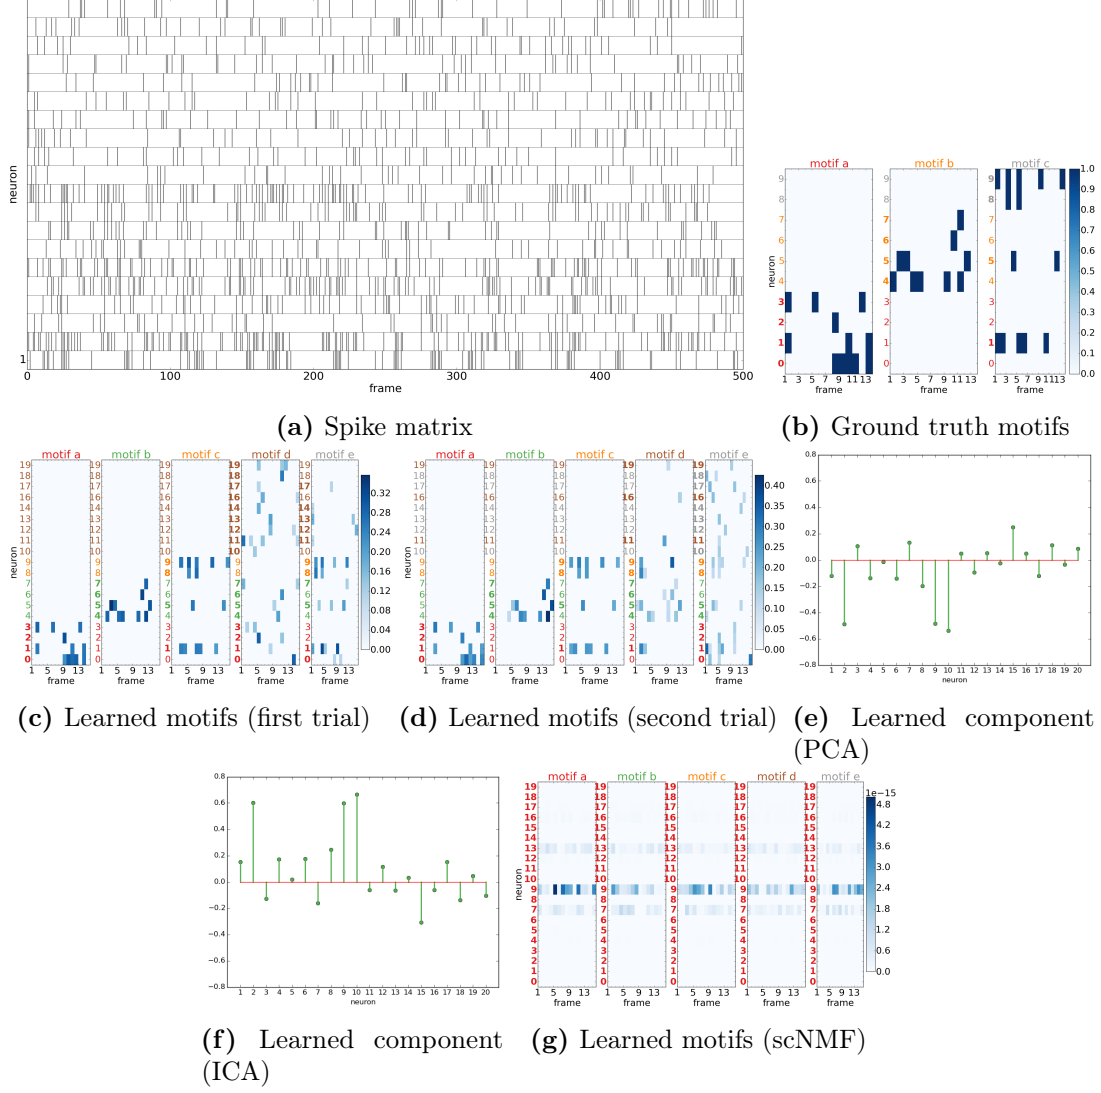


Figure 4.1: Results on a synthetic dataset: **a** shows a synthetic spike matrix with twenty neurons and 500 frames. Each neuron spikes in fifty randomly distributed bins. **b** shows the three motifs made up by the first ten neurons. The time between two activations of a motif has been modeled as Poisson distributed with a mean distance of twenty frames. By running our algorithm with two different random initial states the motifs seen in **c** and **d** are learned. **e**, **f** and **g** show that no other method is able to learn a single true motif.

identify associations for short temporal motifs ($\tau = 7$, figure 4.2b) but only we are able to accurately recover most associations in long motifs ($\tau = 21$, figure 4.2c).

In order to study the stability of our results and the effects of the random initialization of the spike trains we use a non parameterized tests similar to the ones used by other methods to estimate the number of ensembles [29, 32]. Starting with the same generated

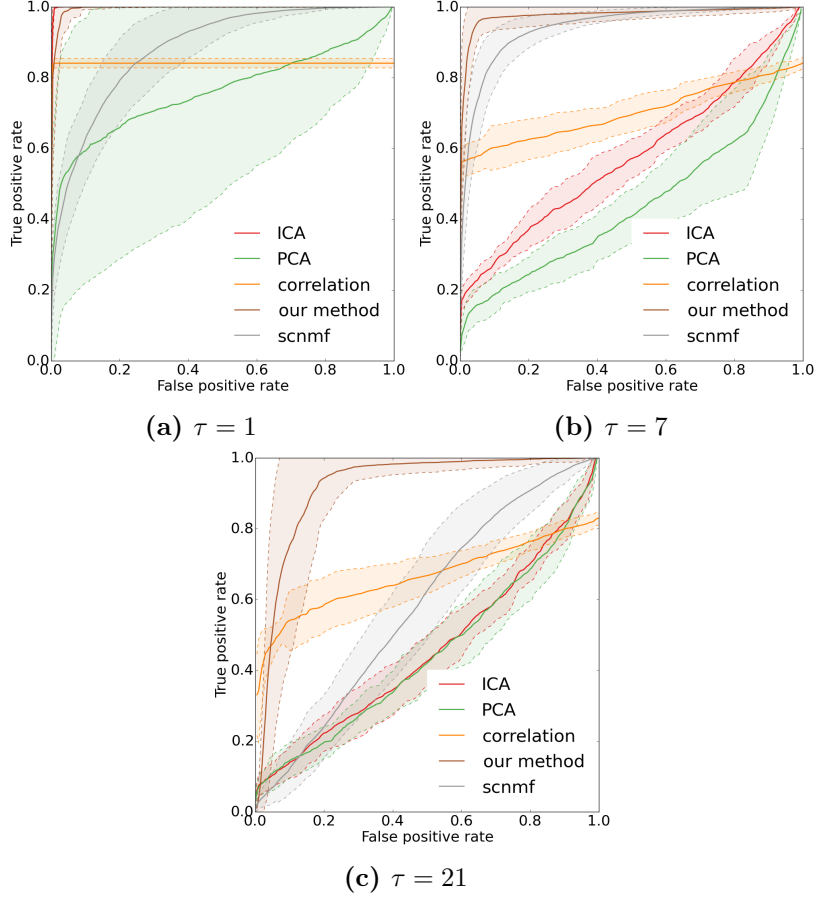


Figure 4.2: ROC curves for different temporal motif lengths: For each of the three experiments twenty different datasets (with different noise levels and amount of neurons shared between ensembles) were generated and all methods were run ten times with different random initialization. We show the mean ROC curve and its standard deviation averaged over all trial done on different synthetic datasets.

dataset consisting of ensembles with no temporal motifs, i.e. subsets of neurons firing synchronously, we shuffle the spike trains of each neuron independently in order to preserve the spike distribution while destroying correlations between individual neurons. We then run our method ten times on the original and the shuffled dataset with different random initializations and compute the average neuron association matrix. Then all unique values inside both matrices are sorted and plotted in figure 4.3d. The point after a significant change in slope in the shuffled matrix is chosen as a threshold on the original matrix. Only five connections are discarded using this procedure.

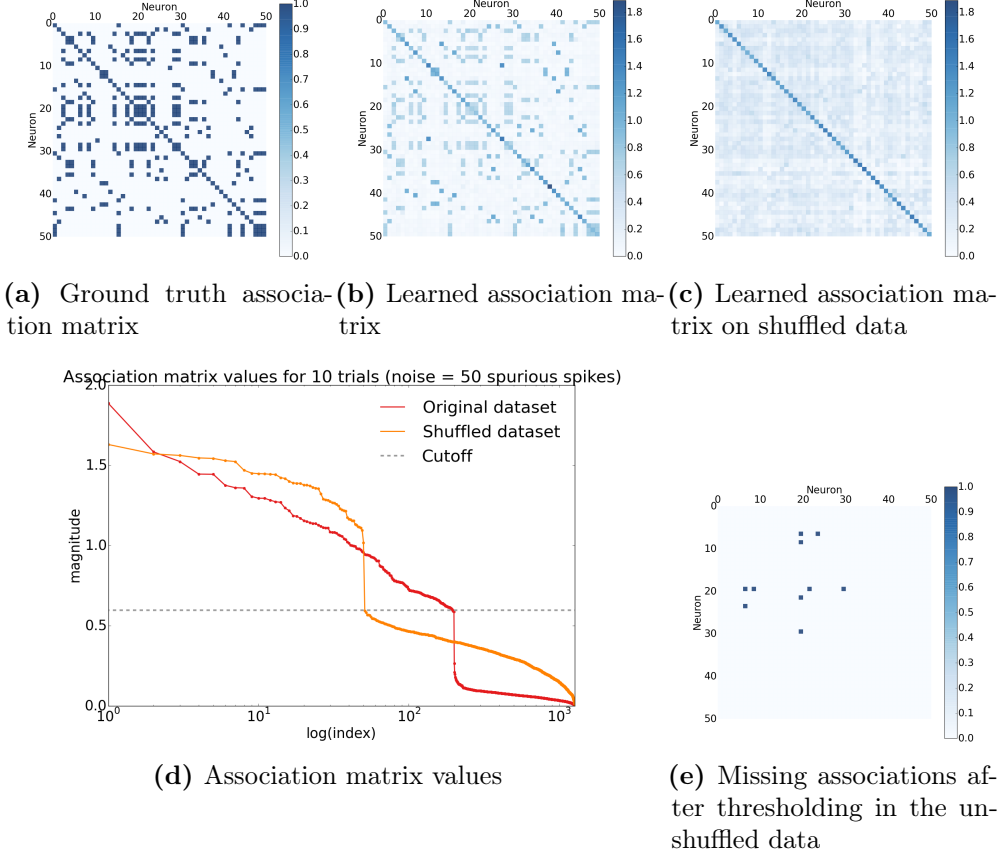


Figure 4.3: Stability analysis of the proposed method: A synthetic dataset consisting of ten synchronous motifs with fifty neurons in total was generated. The spike vector for each neuron is shuffled independently to destroy all correlations. Our proposed method is then run with ten different initial states on both datasets and the resulting average association matrices are computed. A threshold corresponding to a significant change in slope of the shuffled association matrix is chosen and applied to the unshuffled matrix. Using this method only five neuron-neuron associations are not identified.

5 Conclusion

Especially in the past two decades where the number of neurons that can be recorded simultaneously was drastically increased [13, 24, 25] many studies have been done to identify neurons firing in patterns to form motifs as originally suggested by Hebb [1]. While many algorithms have been developed almost all of those are limited to the identification of synchronous motifs only [11, 29, 26, 30, 32].

We have presented a new mathematical method for the identification of motifs that is not limited to synchronous activity. Our method leverages sparsity constraints on the activity and the motifs themselves to allow a simple and elegant formulation that is able to also learn motifs with temporal structure. The proposed algorithm extends convolutional

coding approaches, which have previously already been successful in recovering spike trains from calcium fluorescence recordings [31, 28, 27, 22] and the identification of repeating patterns in audio data [16], with a novel optimization approach to allow modeling of interactions between neurons.

Results on simulated datasets show that the proposed method outperforms others especially when identifying long motifs (figure 4.2).

We hope that these contributions allow to study more complex neuronal firing patterns and help to further the understanding of functional ensembles within the brain as originally suggested by Hebb.

Acknowledgments

This work is a project of the SFB 1134 “Functional Ensembles” funded by the Deutsche Forschungsgemeinschaft (DFG).

References

- [1] D. Hebb. *The Organization of Behaviour: A Neuropsychological Theory*. New York, USA: Wiley, 1949.
- [2] Vladimir Alexandrovich Marchenko and Leonid Andreevich Pastur. “Distribution of eigenvalues for some sets of random matrices”. In: *Matematicheskii Sbornik* 114.4 (1967), pp. 507–536. DOI: [10.1070/SM1967v001n04ABEH001994](https://doi.org/10.1070/SM1967v001n04ABEH001994).
- [3] David Marr, David Willshaw, and Bruce McNaughton. *Simple memory: a theory for archicortex*. Springer, 1991. DOI: [10.1007/978-1-4684-6775-8_5](https://doi.org/10.1007/978-1-4684-6775-8_5).
- [4] Stephane Mallat and Zhifeng Zhang. “Matching Pursuit With Time-Frequency Dictionaries”. In: *IEEE Transactions on Signal Processing* 41 (1993), pp. 3397–3415. DOI: [10.1109/78.258082](https://doi.org/10.1109/78.258082).
- [5] Wolf Singer. “Synchronization of cortical activity and its putative role in information processing and learning”. In: *Annual review of physiology* 55.1 (1993), pp. 349–374. DOI: [10.1146/annurev.ph.55.030193.002025](https://doi.org/10.1146/annurev.ph.55.030193.002025).
- [6] Pierre Comon. “Independent component analysis, a new concept?” In: *Signal processing* 36.3 (1994), pp. 287–314.
- [7] Miguel A Nicolelis, Luiz A Baccala, RC Lin, and John K Chapin. “Sensorimotor encoding by synchronous neural ensemble activity at multiple levels of the somatosensory system”. In: *Science* 268.5215 (1995), pp. 1353–1358. DOI: [10.1126/science.7761855](https://doi.org/10.1126/science.7761855).
- [8] Robert Tibshirani. “Regression shrinkage and selection via the lasso”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* (1996), pp. 267–288.
- [9] Miguel AL Nicolelis, Erika E Fanselow, and Asif A Ghazanfar. “Hebb’s dream: the resurgence of cell assemblies”. In: *Neuron* 19.2 (1997), pp. 219–221. DOI: [10.1016/S0896-6273\(00\)80932-0](https://doi.org/10.1016/S0896-6273(00)80932-0).

- [10] Dimitri P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1999.
- [11] John K Chapin and Miguel AL Nicolelis. “Principal component analysis of neuronal ensemble activity reveals multidimensional somatosensory representations”. In: *Journal of neuroscience methods* 94.1 (1999), pp. 121–140. DOI: [10.1016/S0165-0270\(99\)00130-2](https://doi.org/10.1016/S0165-0270(99)00130-2).
- [12] Per Christian Hansen. “Deconvolution and regularization with Toeplitz matrices”. In: *Numerical Algorithms* 29.4 (2002), pp. 323–378. DOI: [10.1023/A:1015222829062](https://doi.org/10.1023/A:1015222829062).
- [13] György Buzsáki. “Large-scale recording of neuronal ensembles”. In: *Nature neuroscience* 7.5 (2004), pp. 446–451. DOI: [10.1038/nn1233](https://doi.org/10.1038/nn1233).
- [14] Yuji Ikegaya, Gloster Aaron, Rosa Cossart, Dmitriy Aronov, Ilan Lampl, David Ferster, and Rafael Yuste. “Synfire chains and cortical songs: temporal modules of cortical activity”. In: *Science* 304.5670 (2004), pp. 559–564. DOI: [10.1126/science.1093173](https://doi.org/10.1126/science.1093173).
- [15] Hui Zou and Trevor Hastie. “Regularization and variable selection via the Elastic Net”. In: *Journal of the Royal Statistical Society, Series B* 67 (2005), pp. 301–320. DOI: [10.1111/j.1467-9868.2005.00503.x](https://doi.org/10.1111/j.1467-9868.2005.00503.x).
- [16] Paul D O’Grady and Barak A Pearlmutter. “Convolutional non-negative matrix factorisation with a sparseness constraint”. In: *Machine Learning for Signal Processing, 2006. Proceedings of the 2006 16th IEEE Signal Processing Society Workshop on*. IEEE, 2006, pp. 427–432. DOI: [10.1109/MLSP.2006.275588](https://doi.org/10.1109/MLSP.2006.275588).
- [17] Alik Mokeichev, Michael Okun, Omri Barak, Yonatan Katz, Ohad Ben-Shahar, and Ilan Lampl. “Stochastic emergence of repeating cortical motifs in spontaneous membrane potential fluctuations in vivo”. In: *Neuron* 53.3 (2007), pp. 413–425. DOI: [10.1016/j.neuron.2007.01.017](https://doi.org/10.1016/j.neuron.2007.01.017).
- [18] Eva Pastalkova, Vladimir Itskov, Asohan Amarasingham, and György Buzsáki. “Internally generated cell assembly sequences in the rat hippocampus”. In: *Science* 321.5894 (2008), pp. 1322–1327. DOI: [10.1126/science.1159775](https://doi.org/10.1126/science.1159775).
- [19] M. Protter and M. Elad. “Image Sequence Denoising via Sparse and Redundant Representations”. In: *IEEE Transactions on Image Processing* 18.1 (2009). DOI: [10.1109/TIP.2008.2008065](https://doi.org/10.1109/TIP.2008.2008065).
- [20] J.-C. Delvenne, S. N. Yaliraki, and M. Barahona. “Stability of graph communities across time scales”. In: *Proceedings of the National Academy of Sciences* 107.29 (2010), pp. 12755–12760. DOI: [10.1073/pnas.0903215107](https://doi.org/10.1073/pnas.0903215107). eprint: <http://www.pnas.org/content/107/29/12755.full.pdf>. URL: <http://www.pnas.org/content/107/29/12755.abstract>.
- [21] Arthur Szlam, Koray Kavukcuoglu, and Yann LeCun. “Convolutional Matching Pursuit and Dictionary Training”. In: *Computer Research Repository (arXiv)* (2010). URL: <http://arxiv.org/abs/1010.0422>.

- [22] Joshua T. Vogelstein, Adam M. Packer, Tim A. Machado, Tanya Sippy, Baktash Babadi, Rafael Yuste, and Liam Paninski. “Fast non-negative deconvolution for spike train inference from population calcium imaging”. In: *Journal of Neurophysiology* (2010). DOI: [10.1152/jn.01073.2009](https://doi.org/10.1152/jn.01073.2009).
- [23] Vitor Lopes-dos Santos, Sergio Conde-Ocazonez, Miguel AL Nicolelis, Sidarta T Ribeiro, and Adriano BL Tort. “Neuronal assembly detection and cell membership specification by principal component analysis”. In: (2011). DOI: [10.1371/journal.pone.0020996](https://doi.org/10.1371/journal.pone.0020996).
- [24] Ian H Stevenson and Konrad P Kording. “How advances in neural recording affect data analysis”. In: *Nature neuroscience* 14.2 (2011), pp. 139–142. DOI: [10.1038/nn.2731](https://doi.org/10.1038/nn.2731).
- [25] Misha B Ahrens, Michael B Orger, Drew N Robson, Jennifer M Li, and Philipp J Keller. “Whole-brain functional imaging at cellular resolution using light-sheet microscopy”. In: *Nature methods* 10.5 (2013), pp. 413–420. DOI: [10.1038/nmeth.2434](https://doi.org/10.1038/nmeth.2434).
- [26] Ferran Diego and Fred A Hamprecht. “Learning Multi-level Sparse Representations”. In: *NIPS*. 2013. URL: <http://papers.nips.cc/paper/5076-learning-multi-level-sparse-representations>.
- [27] Eftychios A Pnevmatikakis and Liam Paninski. “Sparse nonnegative deconvolution for compressive calcium imaging: algorithms and phase transitions”. In: *NIPS*. 2013. URL: <http://papers.nips.cc/paper/4996-sparse-nonnegative-deconvolution-for-compressive-calcium-imaging-algorithms-and-phase-transitions.pdf>.
- [28] Eftychios A Pnevmatikakis, Timothy A Machado, Logan Grose, Ben Poole, Joshua T Vogelstein, and Liam Paninski. “Rank-penalized nonnegative spatiotemporal deconvolution and demixing of calcium imaging data”. In: *Computational and Systems Neuroscience (Cosyne) 2013*. 2013. URL: <http://www.stat.columbia.edu/~liam/research/abstracts/cosyne-13/eftychios.pdf>.
- [29] Vitor Lopes-dos Santos, Sidarta Ribeiro, and Adriano BL Tort. “Detecting cell assemblies in large neuronal populations”. In: *Journal of neuroscience methods* 220.2 (2013), pp. 149–166. DOI: [10.1016/j.jneumeth.2013.04.010](https://doi.org/10.1016/j.jneumeth.2013.04.010).
- [30] Yazan N. Billeh, Michael T. Schaub, Costas A. Anastassiou, Mauricio Barahona, and Christof Koch. “Revealing cell assemblies at multiple levels of granularity”. In: *Journal of Neuroscience Methods* 236 (2014), pp. 92–106. ISSN: 0165-0270. DOI: [10.1016/j.jneumeth.2014.08.011](https://doi.org/10.1016/j.jneumeth.2014.08.011).
- [31] Ferran Diego and Fred A Hamprecht. “Sparse space-time deconvolution for Calcium image analysis”. In: *Advances in Neural Information Processing Systems*. 2014, pp. 64–72. URL: <http://papers.nips.cc/paper/5342-sparse-space-time-deconvolution-for-calcium-image-analysis>.

- [32] Luis Carrillo-Reid, Jae-eun Kang Miller, Jordan P Hamm, Jesse Jackson, and Rafael Yuste. “Endogenous Sequential Cortical Activity Evoked by Visual Stimuli”. In: *The Journal of Neuroscience* 35.23 (2015), pp. 8813–8828. DOI: [10.1523/JNEUROSCI.5214-14.2015](https://doi.org/10.1523/JNEUROSCI.5214-14.2015). URL: <http://www.jneurosci.org/content/35/23/8813.short>.